



Speed and Scalability of Key Materials Science Code (CASTEP) Quadrupled by HECToR dCSE Team

*Keith Refson, STFC Rutherford Appleton Laboratory
Phil Hasnip, University of York
HECToR CSE Team, Numerical Algorithms Group Ltd (NAG)*

A HPC developer from the University of York, working under NAG's Computational Science and Engineering (CSE) support service for HECToR, the UK's national academic supercomputing facility, has optimized the CASTEP materials science code resulting dramatic improvements in code performance and scalability which could potentially save millions of pounds and allow significant new science to be undertaken for the UK Car-Parrinello Consortium (UKCP).

CASTEP user and Principal Investigator on the project, Dr Keith Refson (STFC) said *"The performance and scaling gains achieved by the band-parallel CASTEP represent a very substantial advance in the efficiency of utilisation of CPU cycles on HECToR. This will not only result in a lower time and cost and more rapid turnaround for jobs already planned, but as intended will permit larger and more complex simulations, using more processors, which were not previously feasible."*

The UKCP Chairman, Dr Matt Probert of the Department of Physics at the University York predicts that due to the speed and scaling gains resulting from CASTEP's improvements the software can now be utilized for larger scale scientific work, i.e. bigger atomic systems for less wall-clock time and/or more simulations– within their existing budgets. In fact, several research projects have been waiting in anticipation of the planned release of the improved CASTEP in 2009.

HECToR

HECToR is managed by EPSRC on behalf of the participating Research Councils with a mission to support capability science and engineering in UK academia. The Cray XT supercomputers, located at the University of Edinburgh, are managed by UoE HPCx Ltd. The CSE Support Service is provided by NAG Ltd and ensures users have access to appropriate HPC expertise to effectively exploit advanced supercomputers for their science. A critical feature of the CSE Support Service is the distributed CSE (dCSE) programme which, through lightweight peer review, delivers dedicated performance and scalability projects on specific codes in response to proposals from users. The dCSE programme now consists of nearly 40 focused projects complementing the traditional HPC user applications support and training also provided by NAG.

The dCSE projects completed so far have delivered outstanding examples of the cost savings and new science that can be enabled through dedicated CSE effort. The CASTEP project reported here was one of the first of these success stories.

Project Background

The objective of the dCSE project was to develop an improved, more scalable version of CASTEP – a commercial and academic software package which uses density functional theory with a plane wave basis set to calculate electronic properties of solids from first principles. The key task of the project was to implement band-parallelism in order to allow the code to scale to more than 1000 cores on HECToR. CASTEP is used on HECToR to model a range of materials or molecules at the atomic level. In particular scientists run CASTEP to obtain information about total energies, forces and stresses on an atomic system, as well as calculating optimum geometries, band structures, optical spectra, phonon spectra as well as molecular dynamics simulations. Dr Keith Refson from the Computational Materials Science Group at the Rutherford Appleton Laboratory, was the Principal Investigator on the project and NAG contracted Dr Phil Hasnip of the Department of Physics at the University of York to carry out the code development work in collaboration with the CASTEP developers and the NAG CSE team.

CASTEP

CASTEP is a commercial (and academic) software package which uses density functional theory with a plane wave basis set to calculate electronic properties of solids from first principles. CASTEP is a fully featured first principles code and as such its capabilities are numerous. Aiming to calculate any physical property of the system from first principles, the basic quantity is the total energy from which many other quantities are derived. <http://www.castep.org/>

Project Results

The results of this work were excellent. The improved code has a speed-up factor of between 2 and 4 times the original and now scales to over 1000 cores against 256 previously.

The UKCP Chairman, Dr Matt Probert of the Department of Physics at the University York estimated that the CASTEP consortium was using around 10m Allocation Units (AU)s per annum on HECToR – a nominal cost of around £640k. Making the code 2-4 times more efficient could result in a saving of £320k-£480k per annum (a saving of around say £1.6m-£2.4m over the remaining life of HECToR); all for around 8 person months of effort. Commenting on the massive return on investment, Dr Probert said *"I guess it goes to show the value of centrally supporting key software packages, and that there is a considerable saving to be made due to scale of usage. Also the HECToR dCSE scheme is well worth supporting and continuing - and that the dCSE postdoc (Phil Hasnip in this case) was very good value for money!"*

A full technical report can be found at <http://www.hector.ac.uk/cse/distributedcse/reports/>

For more information contact: HECToR CSE Team

The Numerical Algorithms Group Ltd, Wilkinson House, Jordan Hill Road, Oxford, OX2 8DR, United Kingdom

Telephone: 01865 511 245

Email: hector-cse@nag.co.uk

Web: <http://www.hector.ac.uk/cse/>

HECToR – A Research Councils UK High End Computing Service